

SCIENCE AND ETHICS: TOWARD A THEORY OF ETHICAL VALUE

John R. Welch

Zeitschrift für Allgemeine Wissenschaftstheorie, 25 (1994), 279–292

ABSTRACT

This paper sketches descriptive and normative components of a theory of ethical value. The normative component, which receives the lion's share of attention, is developed by adapting Laudan's levels of scientific discourse. The resulting levels of ethical discourse can be critically addressed through the use of inductive inference, falsification, and causal inference. These techniques are likewise appropriate to the corresponding levels of scientific discourse.

1. Introduction

The word "value" is a workhorse, outside ethics as well as in. Not only is it much used, however; it is variously used. Within ethics alone we find it as a verb ("to value"), a concrete noun ("the value of an honest report"), and an abstract noun or adjective ("theory of value", "value terms"). Accounting for these uses is a formidable task for any theory of ethical value. Likewise for the uses of ethical value terms themselves ("just," "courageous," "best," "worst," "the highest good," and so on). A theory that could clarify our usage of these terms would not only be descriptive; it would draw the lines, at least in principle, between legitimate and illegitimate uses of the terms in its domain. The theory, that is, would be critical.

To draft such a theory--a critical theory of ethical value-- is the object of this study. I will argue for a use of reason in ethics that parallels the use of reason in the sciences. I will not claim, however, that ethics should somehow be made scientific; we have every reason not to confuse the two. The claim is that scientific reason and ethical reason, by virtue of being reason, have a great deal in common. Just what that is I hope to flesh out in the following pages.

Part 2 establishes some working distinctions to be used in articulating the theory. Part 3 identifies, with the help of some recent work in the philosophy of science, three levels of ethical discourse. These levels become the scaffolding for a normative theory of value, whose contours emerge in Part 4.

Before proceeding, I want to say something about what the theory articulated here is not. It is not ontological in the sense that it tries to determine what ethical values really are. The theory, rather, is linguistic: though we use these terms all the time, how do we make sense of them? Moreover, the theory falls well short of a full-fledged theory of value. It is not literally a theory of value, first of all, for nonethical values (scientific, political, aesthetic, legal, religious, logical, medical, ludic, etc.) are not actively considered. And it is certainly not a full-fledged

theory; it concentrates on overall structure, leaving many matters of detail to be filled in. The idea is to present the big picture in a few strokes.

2. Some Preliminary Distinctions

Let us take sentences like the following as representative of the uses to be accounted for by the theory.

- (1) She values friendship above all else.
- (2) The present policy is unfair.
- (3) That shows the value of courage.
- (4) Be honest so that people will respect you.
- (5) The highest good is the greatest happiness for the greatest number.
- (6) Never be unjust.

From this small sample, we can draw three preliminary distinctions.

The first is the distinction between values as means and values as ends. (3) and (4) point out, though in very different ways, that there is something that courage and honesty are good for; their value, that is, is as means to ends. But (1), (5), and (6), again in distinctive ways, handle friendship, the greatest happiness for the greatest number, and injustice not as means but as possible goals or ends. Friendship, for the person described in (1), is the supreme end; that person is mistaken, according to (5), for the supreme end is the greatest happiness for the greatest number; and injustice, according to (6), should not be an end at all. Observe that the terms of this distinction are not mutually exclusive, since something can be valued as both a means and an end. Take the money for which one works as an end but which is also a means for feeding one's family, for instance. Or the plumber who is both a means for fixing a broken pipe and a Kantian end in oneself. By contrast, though, the terms of the two working distinctions that follow are mutually exclusive.

To distinguish between value judgment and value direction, let us key on what speech act theorists call the illocutionary act performed by the speaker of the sentence. Value judgments simply affirm that something does or does not possess some evaluative property. (2) asserts that the present policy does not have the property of fairness; (3), that something shows courage has value as a means; and (5), that the referents of the definite descriptions "the highest good" and "the greatest happiness for the greatest number" are identical. Value directions, on the other hand, prescribe or proscribe what one's values ought or ought not to be. (4) prescribes honesty as a means to respect, while (6) proscribes injustice as an end.

Finally, there is clearly a difference between subjective and objective claims using value terms. (1) is subjective in that it does not claim that the value judgment of friendship is justified; it says only that someone thinks it is. (1) can be true, therefore, even if the value judgment it describes is mistaken. The value judgments expressed by (2), (3), and (5), however, are objective in their claims that what is being asserted of a policy, courage, and the highest good are correct. If these claims should turn out to be incorrect, then so of course would (2), (3), and (5). The value directions of (4) and (6) also make objective claims, though not in the same way. They are not statements of fact like "Someone says 'Do such and such,'" which would make them value-directive analogues of (1). Rather, they claim to bind the persons to whom they are addressed: "Do such and such." They assume directive authority.

This final working distinction permits us to delimit a normative and a descriptive component of the theory of ethical value. The descriptive component deals with subjective value claims like (1), whereas the normative component concerns objective claims like (2)-(6). Though the descriptive theory of value has its share of deep problems, the normative theory is the more philosophically problematic of the two. I will concentrate on it in what follows.

3. Levels of Ethical Discourse

Sentences like (2)-(6) are so different that a way of relating them at some deeper level is a near necessity for a normative theory of value. We are fortunate to find one, not quite ready made but easily adaptable, in recent work on the philosophy of science.

Larry Laudan has proposed in Science and Values what he calls the reticulational model of scientific rationality.¹ Perhaps the simplest way to introduce its central ideas is by contrast with the older, hierarchical model of justification associated with Popper, Hempel, and Reichenbach, among others. The proponents of this model identified three levels of scientific discourse. The first is factual, wherein scientists might agree or disagree about the structure of DNA or the existence of phlogiston. The second is methodological; agreement or disagreement here could be over rules of scientific procedure such as the Royal Society's "Nullius in verba" or how to use a bubble chamber. The third, axiological level concerns the goals of science: to avoid action at a distance, say, or to strive for falsifiable theories. How the model explains the resolution of disagreement clearly reveals its hierarchical structure. Disagreements at the factual level are resolved by ascending to the methodological level; disagreements over methodology are resolved by moving up to axiology. But what about axiological differences? That was the flaw, of course. The model provided no resources whatsoever for resolving them in a rational way. And since the model makes all justification depend ultimately on axiology, to claim that axiology is impervious to reason is to admit that, ultimately, so too is science. The hierarchical model builds science on sand.

To avoid that problem, Laudan does away with the hierarchy. The reticulational model maintains the distinction between factual, methodological, and axiological discourse. But it subjects these levels to "a kind of leveling principle that emphasizes the patterns of mutual dependence between these various levels."² Whereas the hierarchical model permitted justification to flow only downward, from axiology to methodology to fact, the reticulational

model allows justification to flow upward as well. Not only, then, are facts justified by methods and goals, but methods and goals are justified by facts.

Laudan is explicit about the possibility of extending the reticulational model to moral theory, though he does not try to do so in his discussion of values in science. But there is nothing to stop us from making a start. Let us observe, to begin with, that we can distinguish levels in ethics analogous to Laudan's levels in science. (2), for example, belongs to the ethically factual or, as I shall say, to the ethically descriptive level; (3) and (4), to ethical methodology; and (5) and (6), to ethical axiology. To draw this analogy is to acknowledge no more than that the surface grammars of science and ethics, each in its peculiar way, concern ends, means, and the bearers of properties. The analysis of these levels then proceeds as it may--along cognitive or noncognitive lines, as one sees fit. Thus claims that are ethically descriptive could be explained in prescriptivist or emotivist terms, though that is not what I propose to do. The analysis offered below is, with some important reservations, cognitive.

4. The Normative Theory of Ethical Value

4.1. Ethical Description

Let us begin at the descriptive level with (2). Notoriously, there are situations in which some people firmly believe (2), while others, equally firmly, believe its contradictory. But it is just as true, and often overlooked, that there are other situations in which (2) is so patently correct that it is not even worth asserting. (Take this situation, for example, and "the present policy" to be the reader's reference to anti-Semitic hiring practices.) (2) is not unique in this; any number of sentences make ethically descriptive claims that are controversial at times, trivial at others. Moreover, a great many nonethical descriptive sentences, even scientific ones, have the same trait. "That creature is an animal" is trivial if the reference is to Bossie the cow but controversial if it is to a specimen of Euglena, which photosynthesizes like plants yet has a

flagellum and an eyespot like protozoa. All these sentences act as if they had a short circuit somewhere--epistemologically, that is, they flicker.

The epistemological problem with these sentences has nothing to do with the ambiguity of terms like "the present policy" and "that creature," which can refer to one thing or another, depending on the context. Such ambiguities, built right into the language, are almost invariably resolved in context. The problem, rather, is with the vagueness of predicates like "unfair" and "animal." Though they are vague in very different degrees, they are alike in that their extensions are partly but not entirely clear. Among their range of possible application, then, are two groups of cases: those to which the predicate clearly applies, and those to which its application is unclear. The sentence flickers, epistemologically speaking, as the predicate is applied first to a case from one group, then to a case from the other.

To point an accusing finger at vagueness is merely to name the problem; doing something about it would require much more. We can make a start in that direction by reflecting generally on sentences that attribute some predicate, ethical or not, to a determinate object. These sentences--call them simple categoricals--have unquantified forms like "The F is G," "This F is G," "That F is G," and "f is G."³ Let us consider four logical contexts in which they arise, arranging them epistemologically from less to more primitive. The first context, where simple categoricals figure as conclusions, is deductive: "All F are G, this is an F, hence this F is G," to take a standard case. Asking where universally quantified premises like "All F are G" come from gives us two more contexts. They could come from an exhaustive enumeration of F: "This F is G, this other F is G, ..." until every F has been examined. Or they could come from an inductive enumeration of a sample of F: "This F is G, this other F is G, ..., hence all F are G." In either case, the premises are simple categoricals. If we ask where these premises come from, the answer is still another context: the form of induction known as eduction, where the inference is from particular to particular. Having seen a furry, bounding thing turn out to be a rabbit before,

we might infer that this furry, bounding thing is a rabbit. Or knowing that this green, crispy leaf is lettuce, we might conclude that that green, crispy leaf is lettuce. Here simple categoricals function as both premises and conclusions. The conclusions come from the premises, of course, and the premises come ultimately from the paradigm cases used in teaching the words: "This is an F," someone says, "and it is G."

As the justificatory relationships among the foregoing examples show, eduction plays a foundational role in epistemology. Classifying ordinary objects as stop signs and porcupines and pears is almost invariably performed eductively, and without that ability it is questionable whether we could know anything at all about the world. Eduction's role at the very base of our knowledge has a lot to do, I believe, with the omnipresence of the epistemological flicker we remarked earlier. The cause of that flicker, I suggested, was vagueness. I now want to show how easy it is for eduction, being the rude tool that it is, to contribute to that vagueness.

Let us imagine some ancient sailor with a paradigm case of a fish in mind. The sailor knows that that aquatic animal--a shark, perhaps--is a fish. Then it would be perfectly natural to reason eductively that since this whale is an aquatic animal, this whale is a fish. Aristotle, on the other hand, thinking of some human being as a paradigm case of a mammal, would make the eductive inference that since that animal that nourishes its fetus with a placenta is a mammal, and since this whale is an animal that nourishes its fetus with a placenta, this whale is a mammal. Put the sailor and Aristotle together, and the result is what I will call convergent eduction: two chains of eductive inference converging at the same point--in this case, the whale. One chain reasonably identifies it as a fish; the other, just as reasonably, as a mammal. Until the work of Linnaeus and Müller definitively settled the matter, the predicates "fish" and "mammal" had billions of clear cases at the same time that the Cetacea were taxing taxonomists' wits. To that extent at least, the predicates were vague.

We find the same kind of problem, and occasionally the same kind of progress, in ethics. The African slave trade was often justified in the language of Aristotle's Politics, where the soul is said to govern the body, a human being an animal, a parent a child, all according to the natural dominance of inferior by superior. Referring to one of these cases, a defender of the slave trade could reason as Aristotle himself would have: since that case of superior governing inferior is just, and since the treatment of Africans in the New World is a case of superior governing inferior, the treatment of Africans is just. But Bartolomé de Las Casas, after having lobbied for the use of African rather than native American slaves in the New World, changed his mind. Conscious of the injustice of the enslavement of the Indians, he reasoned that since that policy of slavery is unjust, and since the treatment of Africans in the New World is a policy of slavery, the treatment of Africans is unjust. Once again, then, we have eductive chains of inference converging at the same point. One says that the slave trade is just, the other that it is unjust. How long it took the rest of the world to catch up with Las Casas on this point is as painful to contemplate as the lag in catching up with Aristotle on the whale. But catch up it did, making the predicate "just" that much clearer than before.

The clash of reasoned opinion that comes from the convergence of eductive chains can be transmitted almost spontaneously up the epistemological ladder. Having concluded through repeated eductions that a number of F are G, we could leap to the inductive conclusion that all F are G, which could serve in turn as a premise in a deductive inference that some as yet unexamined F is G. At the same time, convergent eductive conclusions that various F are not G could ground the inference that all F are not G and, as a result, that some unexamined F is not G. The respective chains of inference just get further and further apart, it seems, the differences more and more unbridgeable. Is there any rational way to bring them together?

There are many ways, actually. Some of them work piecemeal, as the histories of the whale's taxonomy and the antislavery movement show. But there is one that is at once more

sweeping, more promising, and more problematic than any of these specific approaches. It is the task of defining inductive validity. If we could state how the premises of an inductive argument should support its conclusion with anything like the exactness with which it can be done for deduction, neither the empirical sciences nor ethics would ever be the same again. Since both rely so heavily on induction, and since induction is a form of induction, a clear criterion of inductive validity would weed out invalid inductions from the start, preventing them from crowding good seed. Being able to do so would no more permit us to do armchair science or ethics than defining deductive validity does. But it would make those armchairs more productive.

What the conditions for inductive validity are is of course a highly technical matter, and I will have no more to say about that here except for two brief remarks.⁴ Even if these conditions could be stated, we should not expect the problems of convergent induction to just vanish, first of all. Because inductive reasoning is probabilistic by nature, it would still be possible for two valid inductions to converge at the same point. Having the conditions for validity in hand merely cuts this possibility to a minimum. Secondly, whatever these conditions turn out to be, they need to reflect our intuitive sense that arguments built on a wide inductive base are better than those built on a narrow one. The more white things seen, the less chance of confusing virgin wool and snow. Likewise in ethics; the more dangerous actions observed, the less confusion between courage and foolhardiness. Besides the purely logical criteria for validity, then, experience is essential for good judgment at the ethically descriptive level.

4.2. Ethical Methodology

Let us now turn to sentences like (3) and (4), which are proper to ethical methodology. In terms of the working distinctions established at the essay's outset, (3) is a value judgment, while (4) is a value direction. We will discuss each in turn.

Suppose that an observer has witnessed what she takes to be a courageous action, and that she believes this action to have saved a child from being run over by a car. Deeply impressed, the observer asserts (3). Despite its apparent simplicity, (3) is a rather complex sentence. It claims that a certain action is courageous; that a certain event is the saving of a child; that the action caused the event; and that similar causes (other courageous actions) tend to produce similar effects (saving people from harm). If any of these claims is false, the observer is wrong; but she is right if all of them are true.

Verifying one of these claims in particular presents special problems: the claim that courageous actions tend to save people from harm. This claim is undoubtedly either true or false, but which is it? How could we tell? The answer, I think, is in the same way we test statistical generalizations in the social sciences--through careful observation and inductive inference. Claims like the foregoing are empirically no different than Gresham's law, for example, in economics.⁵ What is different is that, while prodigious amounts of energy and ingenuity have been expended on empirical testing in the social sciences, to my knowledge an empirical test of a statistical generalization has never even been attempted in ethics. And that, I am afraid, is a serious indictment of traditional ethics.

Little need be said about verifying the other claims, I think. Whether the courageous action caused the event of saving the child is a straightforward empirical question. And whether the action really is courageous, the event really a saving of the child, is in each case a matter for inductive inference. The inference might be a direct eduction from other cases, or it might be mediated by a universal generalization extrapolated from an inductive base, but in either case the problems and prospects are conditioned by our earlier remarks on convergent eduction and inductive validity.

In discussing (4), it will help to keep in mind the foregoing remarks on methodological value judgments, for it is from them that methodological value directions like (4) are derived.

No doubt the preceding clause will sound outrageous to some, given Hume's exposé of fraudulent derivations of "ought" from "is." But the claim is not that we can derive an ethical "ought" from a nonethical "is", thereby bridging the gap between ethical evaluation and nonethical description. It is that we can derive, in a sense to be clarified below, an ethical "ought" from an ethical "is." Still, someone might reply, whatever we might think of attempts like Searle's to show that, where institutional facts are concerned, we can indeed derive "ought" from "is",⁶ surely Hume was on to something; surely we cannot in general derive sentences of the form "_____ ought to _____" from sentences of the form "_____ is _____," whether they are ethical or not. I agree. Hume was right about that--for standard deductive logic. But from the point of view of erotetic, question and answer logic, he was wrong.

To see why, first take the nonethical case of a master potter who instructs an assistant to fire a freshly enameled pot. The attentive assistant asks why, and the potter answers by stating that firing is the only way to fix the colors. The original instruction was thus a hypothetical imperative: to fix the colors, fire the pot. The relationship between the "ought" of the potter's imperative and the "is" of her judgment is that the latter warrants the former. Consequently, the relationship is analogous to that between the premises and the conclusion of a deductive argument. Just as a challenge to a deductive conclusion must be met by adducing the premises, a challenge to the potter's imperative must be met by stating her judgment. In that analogous sense, then, the "ought" of the imperative is derived from the "is" of the judgment.

Suppose now that a mother has enjoined her child to be honest, and that the child questions the imperative by demanding to know why. The mother responds by stating that honesty is necessary for gaining the respect of others. Her original imperative, then, had the hypothetical structure of (4): Be honest so that people will respect you. The relationship between the "ought" of her imperative and the "is" of her judgment is just as it was in the case of the potter: the judgment warrants the imperative. In the same, analogous sense, therefore, we can

say that the "ought" of the mother's imperative is derived from the "is" of her judgment. The former is justified if the latter is true. But is the latter true? That brings us back around to methodological value judgments and the crying need for empirical research that we observed in connection with (3). To devise reliable empirical indicators for concepts like respect would take some doing, admittedly. But the methodological ingenuity that sociologists have shown in managing concepts as slippery as alienation is an indication of what can be done.⁷

There is another, very general resource for evaluating ethical methodology--so general that I believe it applicable to any cognitive method whatsoever, ethical or not. Its provenance is the critique of foundationalist epistemology mounted by Quine and Goodman, among others, who argued that all justification is ultimately circular. Particular conclusions are justified if drawn from the right methods, but methods are justified if they draw the right conclusions. This is far from being a skeptical position; we can adopt it and still distinguish benign from vicious circularity. But it does deny that anything like Aristotle's noûs or Descartes' clear and distinct ideas or early positivism's observation statements can serve as a foundation for certain knowledge. To adapt this point of view to ethics, John Rawls and later Norman Daniels devised the concept of wide reflective equilibrium.⁸ Daniels defined it as

an attempt to produce coherence in an ordered triple of sets of beliefs held by a particular person, namely, (a) a set of considered moral judgments, (b) a set of moral principles, and (c) a set of relevant background theories.⁹

Each of these three sets of beliefs can be corrected by the others; none is primary. The considered moral judgments, for example, can be revised even though they are regarded as highly probable because made under conditions that are unlikely to produce error.

How could the pursuit of wide reflective equilibrium help to evaluate ethical methods? Well, let us assume that a methodological rule (a moral principle, in Daniels' vocabulary) implies the contradictory of a considered moral judgment and that, faced with the inconsistency of the

principle and the judgment, we consider the latter substantially more probable. We could then conclude that, since the rule implies the contradictory of a truth, and hence a falsehood, the rule is unwarranted or false. The procedure is almost as simple as falsification through modus tollens--almost, but not quite.

The epistemological holism associated with Duhem and Quine would likely force a modification in the evaluation of the offending rule. The holist thesis is that it is not the rule by itself but the rule together with various auxiliary hypotheses that imply the falsehood. What is falsified, therefore, is not the rule but the conjunction of rule and auxiliary hypotheses, making it logically possible that the rule is unexceptionable, that the culprit is one of the auxiliary hypotheses. Very well, then. Suppose now that the complex of methodological rules and auxiliary hypotheses known as rule-utilitarianism implies that Caiaphas' proposal to let one man die for the people is morally correct.¹⁰ Suppose also that our considered moral judgment is that Caiaphas' proposal is wrong, and that in this case we opt for the judgment over the implied contradictory. Then we should infer that the conjunction of rule-utilitarian methods and auxiliary hypotheses is false. Now this falsehood is the result of rule-utilitarian methods being false, or it is not. If it is, we know what to do with the methods. If it is not, the problem stemming from one of the auxiliary hypotheses instead, we have every reason to quarantine rule-utilitarian methods until such time as they can be dissociated from the falsehood.

4.3. Ethical Axiology

Of the three levels of ethical discourse, the axiological level is perhaps the most intractable. This is not peculiar to ethics, for scientific axiology is equally obdurate; we have already noted, in introducing the hierarchical model of justification, the positivists' despair over making the decisions of scientific axiology rational. Reichenbach and Popper, for instance, thought that the axiological dispute between realists and instrumentalists was rationally

irresolvable. This part of the positivist legacy continues to influence the theory of human action in general and ethical action in particular. Some, therefore, even if willing to grant what has been said about ethical description and methodology, would deny that we can make any critical sense out of axiological sentences like (5) and (6). Accept them or not, they would urge, but let us have no illusions about the rationality of the decision. It is an arbitrary act of the will.

That this cannot be the whole truth can be deduced immediately from the role of wide reflective equilibrium in ethical methodology. Recall that wide reflective equilibrium is coherence among our considered moral judgments, moral principles, and relevant background theories, where each is defeasible by the others. The moral principles we discussed in the previous section were obviously not axiological, so let us return for the axiological case. As an easy illustration, take the principle that the highest good is the greatest misery for the greatest number. If the positivist line on axiology were correct, we would have no choice but to admit this principle as a rival to the best of Kant and Mill. But it is eminently clear that such a principle implies any number of contradictories of our considered moral judgments. Faced with a conflict between the principles and the judgments, we naturally favor the judgments. The principle, therefore, is falsified. Unlike the more complex methodological illustration involving rule-utilitarianism, this case is so simple that the complicating factor of auxiliary hypotheses just does not arise. Although we do derive the principle's implications with the help of the rules of logic, we are not about to fault them here. Quick as *modus tollens*, then, the blame falls squarely on the principle.

To find a more challenging principle, we might take (5) and prime the discussion by noting several techniques for goal evaluation. Two of them come from Laudan, who proposes them for rational choice in scientific axiology.¹¹ One of them is to show that a goal is utopian, that we have no reason to think that it can really be achieved. This could be done in several ways: by showing that, given our understanding of logic or the laws of nature, a goal is

demonstrably utopian; by showing that, because a goal cannot be described precisely or unambiguously, it is semantically utopian; or by showing that, because there is no criterion for determining if a goal has been achieved or not, it is epistemically utopian. The second technique is, in effect, to harp on the proverb "Practice what you preach," pointing out any inconsistencies that might exist between explicit goals that are acknowledged but not acted upon and implicit goals that are acted upon but not acknowledged. Both these techniques can be adapted for use in ethics, although I will not do so systematically here. Rather, I will limit myself to what is useful for our present concerns.

An additional technique for goal evaluation is available to all but extreme deontologists who grant no ethical importance whatsoever to the consequences of actions. The technique is to milk the social sciences and humanities for information on the consequences of pursuing diverse ethical goals. Synchronically, we might want to focus on the consequences of ethical goals embedded in contemporary cultures. Comparing the egalitarian bent of communism, say, to the libertarian leanings of capitalism could be immensely important for a more moral future. If the approach is diachronic, on the other hand, the consequences of ethical goals held by earlier cultures could be compared with our own. Chronological distance between cultures need not be an obstacle; T.S. Eliot once responded to the objection that the dead are remote from us because we know so much more than they did by observing "Precisely, and they are that which we know."¹²

With these techniques for goal evaluation in mind, let us move on to the value judgment expressed by (5). Whatever else we may want to say about it, (5) is not epistemically utopian in one important sense. Mill called the suspicion that utility does not imply justice "the only real difficulty in the utilitarian theory of morals."¹³ He was admitting, in effect, that if satisfying the greatest happiness principle does not always satisfy justice, (5) would be false. Well then, could we find a case where the greatest happiness principle dictates one thing and justice another?

There seem to be any number of such cases, actually, and at least two ways of generating them. One is to describe a significant utility gain for a majority at the expense of a smaller utility loss for a minority. An example would be Caiaphas' aforementioned proposal that one man die for the people. The other approach is to turn the internal complexity of the greatest happiness principle against it. Assume, for example, that there are two schemes of utility distribution for persons A, B, C, and D. According to scheme I, A gets 0 units of utility, B and C each get 1, and D gets 7. According to scheme II, A, B, C, and D each get 2 units. Which distribution scheme does the greatest happiness principle sanction? It is not clear, as Nicholas Rescher points out, for the principle is a "two-factor criterion"; the greatest happiness is one factor, the greatest number another.¹⁴ If we emphasize the former factor, the principle picks out I, but if we emphasize the latter, it picks out II. But justice, by contrast, clearly picks out II. Hence we could argue as follows. If utility implies justice, the greatest happiness principle sanctions II. But the greatest happiness principle does not sanction II (that is, what it sanctions is not clear). So utility does not imply justice. And if that is correct, (5) is not.

Though the foregoing remarks have been made at utilitarianism's expense, the techniques for goal evaluation are widely applicable. Take the deontologism defended by Kant, for example, in which the consequences of an action carry no moral weight. Mill observed long ago that despite Kant's extreme anti-consequentialist views, he nevertheless took consequences into account in applying the categorical imperative in the formulation of universal law.¹⁵ What Mill was doing, then, was mounting a Laudan-like critique of the inconsistency between Kant's theory and practice, pointing out that, in this instance, Kant failed to practice what he preached.

The only remaining sentence to be discussed is (6). For the sake of comparison, consider an admonition to a beginning basketball player who has just made a mistake: "Never shoot at the other team's basket." Why should the beginner take heed? Because, when playing basketball, shooting at the other team's basket is not a goal. (6) works the same way. Why should we never

be unjust? Because, when living morally, injustice is not a goal. Note that the situation is strongly analogous to (4), a value direction warranted by a value judgment. Here the value direction of (6) is warranted by the value judgment that injustice is not a moral goal. Like (4), (6) is derived from a value judgment in a sense analogous to that in which we say that a warranted conclusion is derived from its warranting premises. But how do we know if the value judgment is true? Well, how do we know that shooting at the other team's basket is not a goal in basketball? Because players generally shoot at their own basket, they are sanctioned when they do not, and it is written in the rule books. Similarly, we know that injustice is not a moral goal because moral people usually try to be just, they are sanctioned when they are not, and it is written in the rule books. Absent conflicting evidence, then, there is every reason to accept it without hesitation.

Though I am proposing an antipositivist line in axiology, I want to be careful not to overstate my case. The argument is not that all axiological choices are rationally decidable; it is, rather, that some of them are. That it would be overly optimistic to make the universal claim can be suggested quickly by returning for a moment to methodology. For any method whatsoever, ethical or not, it is always possible and sometimes actually the case that some other method is equally effective relative to the criteria at hand. We have an emblem for this sort of dilemma: Buridan's ass, desperately hungry but methodologically unsure, unable to decide between equally near, equally fresh, equally succulent piles of hay. So even in methodology, where the problems are generally less intransigent than those of axiology, reason does not yield a unique solution in every case. It would be unrealistic, then, to expect more of it in the profounder strata of axiology. Perhaps that is just as well; the more rational determinacy, the less room for the creative use of reason, for the novelties born of sense and restlessness.

5. Conclusion

The foregoing story has a moral: to be truly critical, a theory of ethical value has to employ two kinds of resources. In the normative theory of value, experience is required at every step of the way. It is needed at the descriptive level to identify instances of ethical kinds; at the methodological level, to link cause and effect; and at the axiological level, to determine what is utopian and what are the consequences of pursuing certain goals. Thus Aristotle was on the right track in denying phronēsis to the young.¹⁶ Level by level, however, there is also a need for the light of logic. For resolving questions of ethical description, we want criteria of inductive as well as deductive validity; for methodological questions, the techniques of falsification; and for axiological questions, falsification plus criteria for inconsistency. To make any headway in ethics, in short, we need to build on both experience and logic.

But as any philosophy of science will attest, these resources are the building blocks of the sciences. We have seen the parallels. The same levels of discourse are operative in science and ethics, and the logical manipulation of experience within each level is comparable as well. Granted that what scientists and ethicists are talking about is different; but how they go about it, when they do it well, is in many ways the same.

NOTES

¹ Larry Laudan, Science and Values (Berkeley: University of California Press, 1984).

² *Ibid.*, p. 63.

³ On the surface, "Francis is a general" is not categorical in the same way as "The fox is gray"; "Gf" is not the same as "The F is G." But if we assimilate proper names to definite descriptions, rewriting "Francis," for example, as "the object that is Francis," "Francis is a general" turns out with the form "The F is G." See W. V. Quine, Methods of Logic, 4th ed. (Cambridge, MA: Harvard University Press, 1982), pp. 278-83. This maneuver is particularly helpful in handling abstract singular terms like "courage," "slavery," "play," and so forth.

⁴ Quantitative inductive methods are described by Jaakko Hintikka in "A Two-Dimensional Continuum of Inductive Methods," Aspects of Inductive Logic, ed. Hintikka and P. Suppes (Amsterdam: North-Holland, 1966), pp. 113-132; and by Hintikka and I. Niiniluoto, "An Axiomatic Foundation for the Logic of Inductive Generalization', in Formal Methods in the Methodology of Empirical Sciences, ed. M. Przełęcki, K. Szaniawski, and R. Wójcicki (Dordrecht: D. Reidel and Wrocław: Ossolineum, 1976), pp. 57-81.

⁵ Cf. Hempel's remarks on instrumental judgments of ethical value in Aspects of Scientific Explanation (New York: The Free Press and London: Collier-Macmillan, 1965), pp. 84-85.

⁶ John R. Searle, Speech Acts (Cambridge: Cambridge University Press, 1969), pp. 175-98.

⁷ Two seminal studies are Melvin Seeman, "On the Meaning of Alienation," American Sociological Review, 24 (1959), 783-91, and Melvin Kohn, "Occupational Structure and Alienation," American Journal of Sociology, 82, No. 1 (1974).

⁸ John Rawls, A Theory of Justice (Cambridge, MA: Belknap Press, 1971), pp. 20, 48-50.

Norman Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics," The Journal of Philosophy, 76 (1979), 256-82.

⁹ *Ibid.*, p. 258.

¹⁰ Alan Donagan is responsible for the suggestion that rule-utilitarianism implies Caiaphas' proposal. See his contribution to Contemporary Utilitarianism, ed. Michael D. Bayles (Garden City, NY: Doubleday, 1968). The proposal is recorded in John 18, 14.

¹¹ Laudan, pp. 50-62.

¹² T. S. Eliot, "Tradition and the Individual Talent," in Selected Essays, 3rd ed. (London: Faber and Faber, 1951), p. 16.

¹³ John Stuart Mill, "Utilitarianism" in Essays on Ethics, Religion and Society, ed. J.M. Robson (Toronto: Univ. of Toronto Press and London: Routledge & Kegan Paul, 1969), p. 259.

¹⁴ Nicholas Rescher, Distributive Justice (Indianapolis: Bobbs-Merrill, 1966), pp. 25-28. Similar distribution schemes appear on p. 25.

¹⁵ "Utilitarianism," p. 207. Mill refers to the applications of the imperative in the Second Section of Kant's Grounding for the Metaphysics of Morals.

¹⁶ Aristotle, Nicomachean Ethics 1142a13-16.